The Case for System Redundancy in Automated Conflict Detection in Aviation: Reducing the False Alert Problem

Presented by
Ellis Feldman (ellis.feldman@faa.gov)
Federal Aviation Administration
ASD-400, Investment Analysis & Operations Research

INFORMS Conference
Miami Beach, Florida
November 5, 2001

Examples of Detection/Warning Systems

- Smoke Detectors
- Conflict Alerting Systems (FAA Prototype – User Request Evaluation Tool – URET)
- Medical Imaging/Diagnostic Testing
- Baggage Screening

Attributes of Automated Warning Systems

- Warning systems vary in their ability to discriminate unsafe conditions from safe conditions as a function of the:
  - Effectiveness of the detection algorithms
  - Accuracy/reliability of the input data

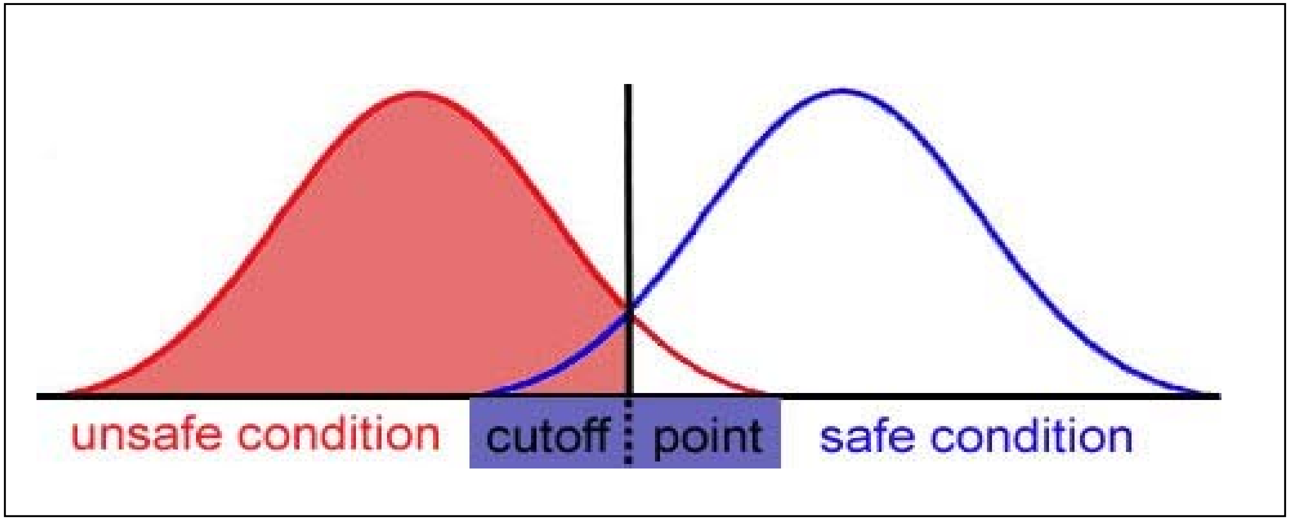

Attributes of Automated Warning Systems (Cont'd)

- Warning systems can be manipulated in terms of the amount of evidence required for an "unsafe" decision
  - Lenient vs. strict alerting threshold
- For a fixed ability to discriminate, increasing the probability of detecting dangerous conditions also increases the probability of false positives.

Definitions

- P_d – Probability of an alert given an unsafe condition
- P_{fa} – Probability of an alert given a safe condition
- p – Prior odds of an unsafe condition

- L – Likelihood ratio = relative likelihood of an alert in the presence vs. the absence of an unsafe condition = $P_d \div P_{fa}$

- PPV – Positive predictive value of an alert = apostiori odds of an unsafe condition given an alert

- Fraction of alerts which are true equals PPV ÷ (1+PPV)

Practical Concerns

- Operators may respond slowly or not at all to warnings if the fraction of alerts which are true is too low, possibly below 80 or 90 percent.

- The PPV sufficient to ensure a reliable response is difficult to attain
  - Dangerous situations are usually rare events
  - Problem worsens as you make the alerting threshold more lenient (in order to get a high probability of detecting a dangerous condition; get more false alerts, for a fixed ability to discriminate)

The Bayesian Approach

- The odds ratio form of Bayes' Theorem:

  Positive Predictive Value = Prior odds · Likelihood ratio

  $$PPV = p \cdot L = p \cdot (P_d \div P_{fa})$$

- Some examples:

Prior odds	P_d	P_{fa}	PPV	Fraction of Alerts Which are True
.001	.99	.05	.0198	2%
.001	.90	.003	0.3	23%
.001	.99	.001	.99	50%
.001	.99	.00011	9*	90%
.0001	.99	.001	.099	9%

* It is only here that the number of true alerts (warnings indicating a dangerous condition) is greater than the number of false alerts.

Mitigations

- Choose an optimal alerting threshold based on prior odds and relative costs and benefits of the two kinds of errors; namely, failing to detect a dangerous condition $(1 - P_d)$ versus falsely reporting danger (P_{fa}). For conflict probe, the threshold may vary depending on the distance between the pair of aircraft.

- Use multiple independent probes

- $L = (P_d^1 \div P_{fa}^1) \cdot (P_d^2 \div P_{fa}^2) \cdot (P_d^3 \div P_{fa}^3) \cdot \ldots \cdot (P_d^N \div P_{fa}^N)$ for N independent probes

Mitigations (Cont’d)

- For the examples shown in the previous table, adding a single backup probe with P_{fa} = .05 and P_d = .99 would increase PPV by a factor of 19.8.

Prior odds	P_d (First Probe)	P_{fa} (First Probe)	PPV (N=2 Probes)	Fraction of Alerts Which are True
.001	.99	.05	0.39	28%
.001	.90	.003	5.9	86%
.001	.99	.001	19.6	95%
.001	.99	.00011	178	99%
.0001	.99	.001	1.96	66%

- System $P_d = .99^N$ if each probe has a 99% P_d.

First Probe P_{fa} Required for Given PPV

- $P_{fa} = (1/PPV) \cdot P_d{}^N \cdot p \div (1 - p) \div Backup_{fa}{}^N$

- P_{fa} for the first probe scales:
  - Linearly in conflict base rate (since p is << 1)
  - Inversely with PPV
  - Log-linearly in *N*
  - Inversely with the *N*th power of the backup false alert rate

Required First Probe P_{fa} (Cont'd)

PPV	Prior odds	N=1	N=2	N=3
9	.001	0.00011	0.0022	0.043
9	.0002	0.000022	0.00044	0.0086
9	.0001	0.000011	0.00022	0.0043

- N is the number of probes
- Each backup probe has a .05 false alert rate
- Each probe has a .99 detection rate: System $P_d = .99^N$
- PPV = 9: the fraction of alerts which are true = 90%.

Conclusions

- Developers of automated conflict warning systems should consider the interactions among the
  - Alerting threshold
  - Prior probability of a dangerous condition
  - System sensitivity
- Multiple independent versions of automated conflict probes may be a better strategy than relying on a single version for reducing the false alert problem in automated conflict detection.
- Whether such independent versions can be implemented remains to be determined.
  - Algorithmic diversity does not guarantee independence
  - Common data inputs (e.g., erroneous forecast winds) could cause multiple versions to miss an alert
- Consider partitioning problem: Apply probes separately to portion of problem where they discriminate best.

References

- Egan, J. (1975). Signal Detection Theory and ROC Analysis. Academic Press, Inc: New York, NY.

- Feldman, E. & Weitzman, D.O. (1999). The false alert problem in automated conflict detection: Crying wolf too often. Journal of Air Traffic Control, April – June, 51-55.

- Parasuraman, R., Hancock, P.A. & Olofinboba (1997). Alarm effectiveness in driver-centered collision warning systems. Ergonomics, 39, 390-399.

- Parasuraman, R. & Hancock, P.A. (1999). Using the signal detection theory and Bayesian analysis to design parameters for automated warning systems. In Automation Technology Performance, pp 63-67 (M.W. Scerbo & M. Mouloua, eds.) L. Erlbaum: Mahwah, NJ.

- Swets, J. (1992). The science of choosing the right decision threshold in high-stakes diagnostics. American Psychologist, 47, 522-532.